msp

Algebraic & Geometric Topology 15 (2015) 1501–1523

Essential surfaces in highly twisted link complements

RYAN BLAIR
DAVID FUTER
MAGGY TOMOVA

We prove that in the complement of a highly twisted link, all closed, essential, meridionally incompressible surfaces must have high genus. The genus bound is proportional to the number of crossings per twist region. A similar result holds for surfaces with meridional boundary: such a surface either has large negative Euler characteristic or is an n–punctured sphere visible in the diagram.

57M25, 57M50

1 Introduction

Links in S^3 are most easily visualized via a projection diagram. However, obtaining topological and geometric information directly from link diagrams has proved to be a difficult task. Historically, alternating links are one of the few classes of links for which this information has been accessible. For instance, links with prime alternating diagrams contain no incompressible tori by work of Menasco [12] and have minimal-genus Seifert surfaces constructible directly from the diagram by work of Crowell [4] and Murasugi [13]. The goal of this paper is to extend results in this vein to diagrams with a high degree of twisting. To state our results, we must define what this means.

A *bigon* in a link diagram $D(K)$ is a disk in the projection plane whose boundary consists of two arcs in the projection of K. Define an equivalence relation on crossings in a diagram in which two crossings are considered equivalent if they are connected by a string of one or more consecutive bigons. Then a *twist region* of a diagram is an equivalence class of crossings. The minimal number of crossings in a twist region of $D(K)$ is called the *height* of D, denoted $h(D)$, and the number of twist regions of $D(K)$ is called the *twist number*, denoted $t(D)$.

The height and twist number of a diagram turn out to be deeply related to the geometric structure of the link it depicts. Lackenby [11] showed that given a prime alternating diagram, the hyperbolic volume of the link complement is bounded both above and below by linear functions of the twist number $t(D)$. Futer, Kalfagianni and Purcell [6]

Published: 19 June 2015 DOI: 10.2140/agt.2015.15.1501

extended these volume estimates to nonalternating diagrams for which $h(D) \geq 7$; that is, diagrams where every twist region contains at least 7 crossings. Additionally, the results of Futer and Purcell [8] imply that when $h(D)$ is large, there is a close connection between the link diagram and any generalized Heegaard decomposition for the exterior of K.

In this paper, we show that $h(D)$ provides a linear lower bound on the genus of essential surfaces in a link complement. Stating our results precisely requires several definitions.

A link diagram is *prime* if every simple closed curve in the projection plane P that meets $D(K)$ transversely in two points in the interior of edges bounds a disk in P that is disjoint from all crossings of the diagram. A diagram is called *twist reduced* if, for every simple closed curve in P that meets $D(K)$ in exactly two crossings, those two crossings belong to the same twist region. (See Figure 1, left.) We will implicitly assume that the diagram $D(K)$ is connected and alternating within each twist region (so the configuration of Figure 1, right, cannot occur). It is easy to verify that every prime link K has a prime, twist-reduced diagram with alternating twist regions. This can be achieved by first applying a maximal number of type II Reidemeister moves that eliminate crossings and then applying flypes to consolidate crossings into a minimal number of twist regions.

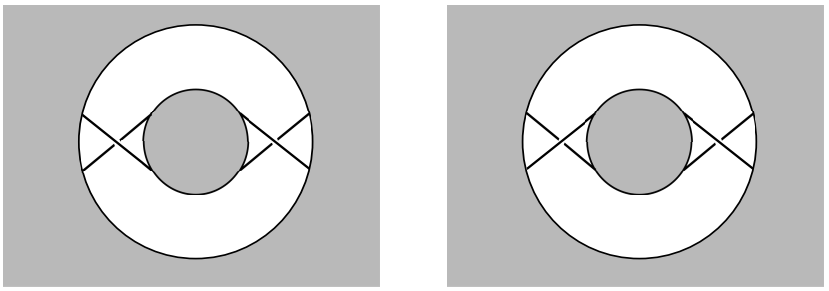

Figure 1: Left: In a twist-reduced diagram, these crossings must belong to the same twist region. Right: In a twist-reduced diagram with alternating twist regions, this configuration cannot occur.

A surface embedded in S^3 is *n–punctured* if it meets K transversely in exactly n points. Two n–punctured surfaces are equivalent if they are transversely isotopic with respect to K. A surface F embedded in S^3 is *c–incompressible* if every disk or 1–punctured disk D embedded in S^3 such that $D \cap F = \partial D$ is transversely isotopic to a disk or 1–punctured disk contained in F while fixing the boundary. Although c–incompressibility is a strictly stronger condition than incompressibility, it is often better behaved than incompressibility and more natural to use when studying surfaces in link exteriors. We can now state the main theorem.

Theorem 1.1 *Let $K \subset S^3$ be a link with a connected, prime, twist-reduced diagram $D(K)$. Suppose $D(K)$ has at least 2 twist regions and $h(D) \geq 6$. Let $F \subset S^3 \smallsetminus K$ be a closed, essential, c–incompressible surface in the link complement. Then we have $\chi(F) \leq 5 - h(D)$.*

Furthermore, if K is a knot, then $\chi(F) \leq 10 - 2h(D)$.

A special case of Theorem 1.1 was proved by Futer and Purcell [7, Theorem 1.4]: if $h(D) \geq 6$, then $\chi(F) < 0$, which implies that F cannot be a sphere or torus.

There is an analogous statement for surfaces with meridional boundary.

Theorem 1.2 *Let $K \subset S^3$ be a link with a connected, prime, twist-reduced diagram $D(K)$. Suppose $D(K)$ has at least 2 twist regions and $h(D) \geq 6$. Let $F \subset S^3 \smallsetminus K$ be a connected, essential, c–incompressible surface in $S^3 \smallsetminus K$ whose boundary consists of meridians of K. Then one of the following holds:*

(i) *F is a sphere with n punctures, which intersects the projection plane in a single closed curve that meets the link n times and is disjoint from all twist regions.*

(ii) *$\chi(F) \leq 5 - h(D)$.*

In other words: either F is "visible in the projection plane", or we obtain the same Euler characteristic estimate as in Theorem 1.1.

There is an interesting analogue between several results involving the height $h(D)$ and results involving the distance of bridge surfaces. The latter is an integer measure of complexity of a bridge surface for a knot that has deep implications for the underlying topology and geometry of the knot exterior. The distance of a bridge surface provides a lower bound on the genus of certain essential surfaces in the knot exterior (see Bachman and Schleimer [2]), while Theorems 1.1 and 1.2 demonstrate an analogous property for height. It is known that both diagrams with large height and bridge surfaces with large distance produce knots with no exceptional surgeries; compare the work of Blair, Campisi, Johnson, Taylor and Tomova [3] to that of Futer and Purcell [7]. Additionally, both height and bridge distance give strong restrictions on the Heegaard surfaces for the knot exterior; compare Tomova [15] to Futer and Purcell [8].

The analogous results about height and bridge distance are all the more striking given that the two notions are in some ways orthogonal. For instance, for 2–bridge knots, distance is essentially equal to the number of twist regions $t(D)$ in a minimal diagram (Zupan [16]), while the height $h(D)$ is the minimal number of crossings per twist region. It would be interesting to know whether the analogous results are indicative of some deeper underlying structure.

Here is a brief outline of the proofs of Theorems 1.1 and 1.2. We begin by adding a number of extra link components to K, so that there is a link component encircling each twist region. (See Figure 2.) In Section 2, we review the construction of this *augmented link* L and show that F can be moved by isotopy into a favorable position with respect to the added link components. In Section 3, we describe a decomposition of the augmented link complement into right-angled ideal polyhedra and again isotope F into a favorable position with respect to these polyhedra.

Sections 4 and 5 constitute the heart of the paper. Here, we use the combinatorics of the ideal polyhedra to estimate the number of times that the surface F must intersect the extra link components that we added to construct L. Each of these intersections will make a definite contribution to the Euler characteristic of F, implying the estimates of Theorems 1.1 and 1.2.

2 Augmented links and crossing disks

In the arguments that follow, we assume that $D(K)$ satisfies the hypotheses of Theorems 1.1 and 1.2. Specifically, $D(K)$ is a connected, prime, twist-reduced diagram with at least 2 twist regions and $h(D) \geq 6$. By Futer and Purcell [7, Theorem 1.4], these hypotheses on $D(K)$ imply that K is prime and $S^3 \smallsetminus K$ is irreducible.

The proof of Theorems 1.1 and 1.2 relies on the geometric study of *augmented links*. Let us recap the relevant definitions, while pointing the reader to Purcell's survey paper [14] for more details.

For every twist region of $D(K)$, we add an extra link component, called a *crossing circle*, that wraps around the two strands of the twist region. The result is a new link J. (See Figure 2.) Now, the manifold $S^3 \smallsetminus J$ is homeomorphic to $S^3 \smallsetminus L$, where L is obtained by removing all full twists (pairs of crossings) from the twist regions of J. This link L is called the *augmented link* corresponding to $D(K)$. By [7, Theorem 2.4], both J and L are prime and $S^3 \smallsetminus L \cong S^3 \smallsetminus J$ is irreducible.

Every crossing circle C_i bounds a *crossing disk* D_i that is punctured twice by strands of K. These twice-punctured disks play a particularly significant role in the hyperbolic geometry of $S^3 \smallsetminus L$. Note that $S^3 \smallsetminus K$ can be recovered from $S^3 \smallsetminus L$ by $1/n_i$ Dehn filling on C_i, where $|n_i|$ is the number of full twists that we removed from the corresponding twist region.

A key goal in proving Theorems 1.1 and 1.2 is to place the surface F into a particularly nice position with respect to the crossing circles and crossing disks. This will be done in two steps. First, we move F by isotopy through $S^3 \smallsetminus K$ into a position that minimizes

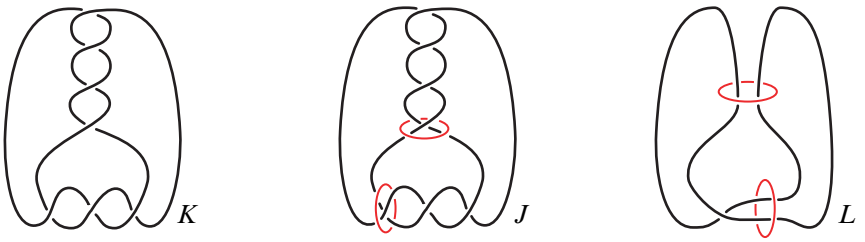

Figure 2: An augmented link L is constructed by adding a *crossing circle* around each twist region of $D(K)$, then removing full twists. The crossing circles are shown in red. Figure borrowed from [7].

the intersections with the crossing disks. Then, in the next section, we drill out the crossing circles and place the remnant surface $F^\circ \subset F$ in normal form with respect to a polyhedral decomposition.

Lemma 2.1 *Let F be a c–incompressible surface in $S^3 \smallsetminus K$ whose boundary (if any) consists of meridians. Move F by isotopy into a position that minimizes the number of components of intersection with the crossing disks. Then every component of intersection between F and a crossing disk D_i is an essential arc in D_i with endpoints in C_i.*

Proof The first step of the proof is to rule out closed curves of intersection. Since D_i is a twice-punctured disk, every closed curve in D_i is either trivial or parallel to one of the boundary components. Isotope F to intersect the union of the D_i minimally.

Since F is incompressible and $S^3 \smallsetminus K$ is irreducible, no curve of intersection can bound a disk in D_i as we could eliminate such a curve of intersection via an isotopy of F. Similarly, since F is c–incompressible and K is prime, no closed curve of intersection can be parallel to a meridian of K. Thus all closed curves of $F \cap D_i$ are parallel to C_i. We may then move F by isotopy in $S^3 \smallsetminus K$ past the crossing circle C_i and remove all remaining curves of intersection $F \cap D_i$, contradicting our minimality assumption. See Figure 3.

Now that we have ruled out closed curves of intersection $F \cap D_i$, all components of intersection must be arcs. An arc with an endpoint on K cannot occur, because F is a meridional surface and after a small perturbation of F we can assume that F is disjoint from the points of intersection between K and the crossing disks. Therefore, every component of $F \cap D_i$ is an arc from C_i to C_i. If any of these arcs are inessential in D_i, then an outermost such arc α can be removed via an isotopy of F supported in a neighborhood of the subdisk of D_i cobounded by α and an arc in C_i. Thus, every component of intersection between F and a crossing disk D_i is an essential arc in D_i with endpoints in C_i. □

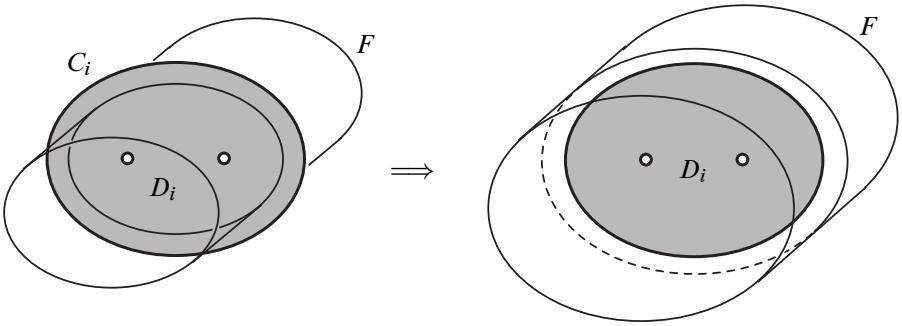

Figure 3: If F intersects a crossing disk D_i in a closed curve, this closed curve must be parallel to the crossing circle C_i and can be removed by isotopy.

Corollary 2.2 *Suppose that F is moved by isotopy into a position that minimizes the number of components of $F \cap \bigcup_i D_i$, as in Lemma 2.1. This position also minimizes the number of points of intersection between F and the crossing circles C_i.*

Proof By Lemma 2.1, every component of $F \cap D_i$ is an arc from C_i back to C_i. Each arc has two endpoints on C_i. Suppose F^* is an isotopic copy of F that minimizes the number of points of intersection between F and the crossing circles C_i. As described in the proof of Lemma 2.1, c–incompressibility of F^* and primeness of K implies we can eliminate loops of intersection between F^* and any D_i that bound disks or 1–punctured disks in D_i via an isotopy that fixes $F^* \cap \bigcup_i C_i$. Similarly, we can remove those loops of intersection between F^* and any D_i that are isotopic to C_i in D_i via an isotopy of F^* that fixes $F^* \cap \bigcup_i C_i$. Hence, we can assume that the points of intersection between F^* and the C_i are in two-to-one correspondence with the components of intersection of F^* and the D_i. Thus, minimizing the number of components of $F \cap \bigcup_i D_i$ also minimizes the number of points of intersection between F and the crossing circles C_i. □

Our next step is to drill out the crossing circles C_i. Suppose, following Corollary 2.2, that F intersects $\bigcup_i D_i$, and thus $\bigcup_i C_i$, minimally. Let $F^\circ = F \smallsetminus \bigcup_i C_i$ be the remnant of F after removing the crossing circles.

Lemma 2.3 *Let L be the augmented link, as in Figure 2. Then, after isotoping F to minimize the number of components of $F \cap \bigcup_i D_i$, $F^\circ = F \smallsetminus \bigcup_i C_i$ is an essential c–incompressible surface in $S^3 \smallsetminus L$.*

Proof Suppose that F° is compressible in $S^3 \smallsetminus L$. Let γ be an essential curve in F° that bounds a compressing disk D in $S^3 \smallsetminus L$. If γ is essential in F, then we contradict

the incompressibility of F. Hence γ bounds a disk E in F that is disjoint from K but meets $\bigcup_i C_i$ in a nontrivial number of points. Since $E \cup D$ is a 2–sphere bounding a 3–ball, there is an isotopy of F taking E to D that strictly reduces the number of components of $F \cap \bigcup_i D_i$, a contradiction.

Suppose that F° is incompressible and c–compressible in $S^3 \smallsetminus L$. Let γ be an essential curve in F° that bounds a 1–punctured disk D in $S^3 \smallsetminus L$. If γ is essential in F, then we contradict the c–incompressibility of F. Hence γ bounds a punctured disk E in F that meets $\bigcup_i C_i$ in at least two points and meets K in at most one point. If E meets K in exactly one point, then, since K is prime, $E \cup D$ is a 2–sphere bounding a 3–ball that meets K in a single unknotted arc. Thus, there is an isotopy of F transverse to K taking E to D that strictly reduces the number of components of $F \cap \bigcup_i D_i$, a contradiction. If E is disjoint from K, then the isotopy of F taking E to D is supported in a 3–ball disjoint from K and again strictly reduces the number of components of $F \cap \bigcup_i D_i$, a contradiction.

Suppose that F° is boundary parallel in the exterior of L. Then either F° is isotopic to the boundary of a regular neighborhood of a component of L, or F° is boundary compressible. If F° is boundary compressible, then, since all boundary components of the exterior of L are tori, either F° is compressible in $S^3 \smallsetminus L$, which is a contradiction as previously demonstrated, or F° is a boundary-parallel annulus. If F° is a boundary-parallel annulus, then F is not essential in $S^3 \smallsetminus K$, a contradiction. If F° is isotopic to the boundary of a regular neighborhood of a component of L, then again F is not essential in $S^3 \smallsetminus K$, a contradiction. □

3 The polyhedral decomposition

In this section, we consider the intersection between the punctured surface F° and a certain polyhedral decomposition of $S^3 \smallsetminus L$. For the purposes of this paper, a *right-angled ideal polyhedron* is a convex polyhedron in hyperbolic 3–space, all of whose vertices lie on the sphere at infinity, and all of whose dihedral angles are $\pi/2$. A *right-angled polyhedral decomposition* of a 3–manifold M is an expression of M as the union of finitely many right-angled ideal polyhedra, glued by isometries along their faces. Note that a right-angled polyhedral decomposition endows M with a complete hyperbolic metric.

In our setting, where M is the augmented link complement $S^3 \smallsetminus L$, there is a well-studied way to decompose M into two identical right-angled ideal polyhedra, first considered by Adams [1] and later popularized by Agol and Thurston [10, Appendix]. Purcell's survey article [14] describes the polyhedral decomposition in great detail.

For our purposes, the salient features are summarized in the following theorem and illustrated in Figure 4.

Theorem 3.1 *Let $D(K)$ be a prime, twist-reduced diagram of a link K with at least 2 twist regions. Let L be the augmented link constructed from $D(K)$. Then the augmented link complement $S^3 \smallsetminus L$ is hyperbolic, and there is a decomposition of $S^3 \smallsetminus L$ into two identical totally geodesic polyhedra P and P'. In addition, these polyhedra have the following properties:*

(i) *The faces of P and P' can be checkerboard colored so that all shaded faces are triangles corresponding to portions of crossing disks and white faces correspond to regions into which L cuts the projection plane.*

(ii) *All ideal vertices are 4–valent.*

(iii) *The dihedral angle at each edge of P and P' is $\pi/2$.*

Proof The hyperbolicity of $S^3 \smallsetminus L$ is a theorem of Adams [1]; cf Futer and Purcell [7, Theorem 2.2]. The remaining assertions are proved in Purcell [14, Proposition 2.2]. □

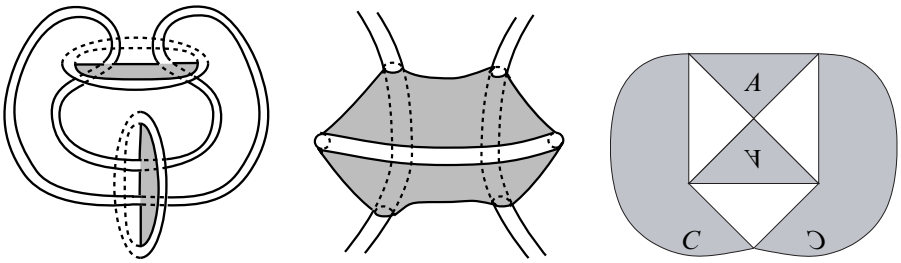

Figure 4: Decomposing $S^3 \smallsetminus L$ into ideal polyhedra. First, slice along the projection plane, then split remaining halves of two-punctured disks. This produces the polygon on the right. Figured borrowed from [7].

Our goal is to place F° in normal form with respect to this polyhedral decomposition. Our convention is that the ideal vertices of the polyhedra are truncated to form *boundary faces* that tile the boundary tori of $S^3 \smallsetminus L$. Then, ∂F intersects the boundary faces in a union of arcs.

Definition 3.2 Let P be a truncated ideal polyhedron. An embedded disk $D \subset P$ is called *normal* if its boundary curve $\gamma = \partial D$ satisfies the following conditions:

(i) γ is transverse to the edges of P.

(ii) γ doesn't lie entirely in a face of P.

(iii) No arc of γ in a face of P has endpoints on the same edge, or on a boundary face and an adjacent edge.

(iv) γ intersects each edge at most once.

(v) γ intersects each boundary face at most once.

If M is a 3–manifold subdivided into ideal polyhedra, a surface S is called *normal* if its intersection with each polyhedron is a disjoint union of normal disks.

It is a well-known fact, originally due to Haken [9], that every essential surface in an irreducible 3–manifold can be isotoped into normal form. However, in our context, we would like to make F° normal while preserving the conclusion of Lemma 2.1. This requires carefully managing the complexity of the surface.

Definition 3.3 Let M be a 3–manifold with a prescribed polyhedral decomposition. Let $S \subset M$ be a properly embedded surface, transverse to the edges and faces of the polyhedra. Order the faces of the polyhedral decomposition as $f_1, \dots, f_n$. Then the *complexity* of S is the ordered n–tuple

$$c(S) = \big(\#(S \cap f_1), \dots, \#(S \cap f_n)\big).$$

Here # denotes the number of components. Given two surfaces S and S', we say that $c(S) \le c(S')$ if the inequality holds in each coordinate. We say that $c(S) < c(S')$ if $c(S) \le c(S')$ and there is a strict inequality in at least one coordinate.

Lemma 3.4 *Let M be an irreducible 3–manifold with incompressible boundary and with a prescribed polyhedral decomposition. Let $S \subset M$ be a properly embedded essential surface, transverse to the edges and faces of the polyhedra. Then S can be isotoped to a normal surface by a sequence of moves that monotonically reduces the complexity $c(S)$.*

The following argument is adapted from Futer and Guéritaud [5, Theorem 2.8], and the figures are also drawn from that paper.

Proof We need to ensure that S satisfies the conditions of Definition 3.2. By hypothesis, S is transverse to the polyhedra. This transversality implies that for every polyhedron P, each component of $S \cap \partial P$ is a simple closed curve, and gives condition (i). Additionally, since S is incompressible, we can assume that we have isotoped S to meet each polyhedron in a collection of properly embedded disks.

Now, whenever some component of $S \cap \partial P$ violates one of the conditions (ii)–(v), we will describe a move that reduces the complexity $c(S)$. That is, for each face σ of the polyhedra, the intersection number $\#(S \cap \sigma)$ will either remain constant or decrease, with a strict decrease for at least one face.

Suppose that γ is a closed curve, violating (ii). Without loss of generality, we may assume that γ is innermost on the face σ. Then γ bounds a disk $D \subset \sigma$ whose interior is disjoint from S. But since S is incompressible, γ also bounds a disk $D' \subset S$. Furthermore, since M is irreducible, the sphere $D \cup_\gamma D'$ must bound a ball. Thus, we may isotope S through this ball, moving D' past D. This isotopy removes the curve γ from the intersection between S and σ. In addition, the isotopy will remove the intersections between D' and any other faces of P.

Next, suppose that γ runs from an edge e back to e, violating the first half of condition (iii). Then γ and a subarc of e cobound a disk $D \subset \sigma$, and we can assume γ is innermost (ie S does not meet D again). We can use this disk D to guide an isotopy of S past the edge e, as in the left panel of Figure 5. This isotopy removes γ from the intersection between S and σ. Some intersection components between S and the interiors of other faces adjacent to e will also merge. Hence, $\#(S \cap \sigma)$ stays constant or decreases for each face.

Suppose that an arc γ runs from a boundary face to an adjacent interior edge in a face σ, violating the second half of condition (iii). Then γ has endpoints in adjacent edges of $\partial\sigma$, and we may assume without loss of generality that it is outermost in σ. Thus γ once again cuts off a disk D from σ. By isotoping S along this disk, as in the right panel of Figure 5, we remove γ from $S \cap \sigma$ and alter the intersection of S with any other face by an isotopy of arcs in that face.

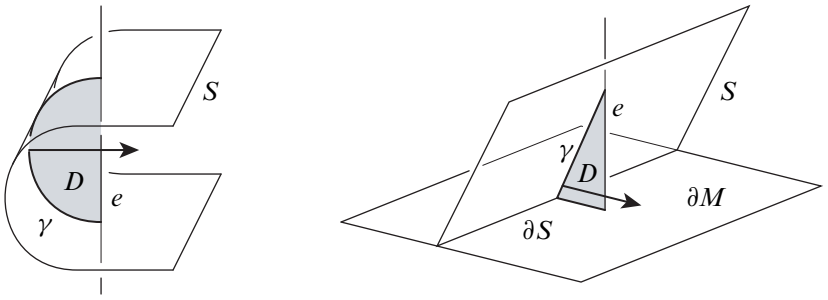

Figure 5: When a surface violates condition (iii) of normality, then an isotopy in the direction of the arrow removes intersections between S and all the faces that meet the edge e.

Suppose a component γ' of $S \cap \partial P$ intersects an edge e twice, violating (iv). Let γ be the closure of a component of $\gamma' - e$ such that γ together with a subarc of e cobound a disk D. By passing to an outermost arc of intersection between S and D, we can assume that $D \cap S = \partial D \cap S = \gamma$. If γ is contained in a face of ∂P, then we violate (iii). Hence, we can assume that γ meets the face of ∂P that contains a neighborhood of $\partial\gamma$ in at least two components. While fixing $\partial D \cap e$, isotope the rest of D slightly into the interior of P. If S meets the interior of D, it does so in simple closed curves. Since S meets P in a collection of properly embedded disks, we can eliminate all components of intersection between S and the interior of D via an isotopy of S that is supported in the interior of P and fixes $c(S)$. After this isotopy, D is a boundary compressing disk for the component of $S \cap P$ that contains γ' in its boundary. As in Figure 5, left, we may use D to guide an isotopy of S past the edge e. Since γ meets the face of ∂P that contains a neighborhood of $\partial\gamma$ in at least two components, this isotopy will strictly reduce $\#(S \cap \sigma)$ for that face and will not increase $\#(S \cap \sigma)$ for any other face σ that meets e.

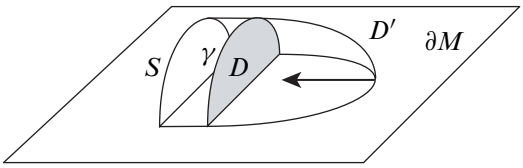

Figure 6: When a surface violates condition (v) of normality, isotoping the disk D' past D removes intersections between S and the faces.

Finally, suppose that γ meets a boundary face twice, violating (v). Then the polyhedron P contains a boundary compression disk D for S such that $\partial D \cap \partial M$ is contained in the boundary face. Since S is boundary incompressible, γ must also cut off a disk $D' \subset S$, as in Figure 6. Since $S^3 \smallsetminus L$ is irreducible, it follows that the disk $D \cup_\gamma D'$ is boundary parallel. Thus we may isotope S through a boundary-parallel ball, moving D' past D, which eliminates all components of intersection between D' and ∂P. Since $\partial D'$ meets the edge of a boundary face, this isotopy strictly lowers $c(S)$.

Since each of the above moves reduces the complexity $c(S)$, a minimum-complexity position will be normal. □

As a consequence, we get the following structural statement.

Lemma 3.5 *Let $F^\circ = F \smallsetminus \bigcup_i C_i$ be as in Lemma 2.3. Suppose that F° has been isotoped into normal form via the procedure of Lemma 3.4. Then the following hold:*

(i) *For each crossing disk D_i, each component of $F^\circ \cap D_i$ is an essential arc in D_i with endpoints in C_i.*

(ii) *For each shaded face σ of the polyhedra, $F^\circ \cap \sigma$ is an arc from an ideal vertex at a crossing circle to the opposite edge. See Figure 7.*

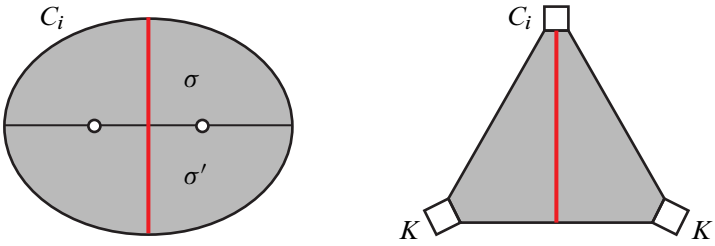

Figure 7: Left: A normal disk D_i is subdivided by the projection plane of L into two shaded faces, one in each polyhedron. By Lemma 3.5, each component of $F^\circ \cap D_i$ is a collection of arcs from C_i to C_i. Right: The picture in a single shaded face σ.

Proof Recall that in the construction of F° we have assumed that we have isotoped F to minimize the number of components of $F \cap \bigcup_i D_i$. Hence, by Lemma 2.1, conclusion (i) holds before we begin the normalization procedure. Additionally, there is an isotopy of F° supported in a neighborhood of the D_i such that after this isotopy any component of $F^\circ \cap D_i$ meets any shaded face of the polyhedra in at most one arc. Since each arc of intersection of $F^\circ \cap D_i$ is essential in D_i, then for each shaded face σ of the polyhedra, $F^\circ \cap \sigma$ is an arc from an ideal vertex at a crossing circle to the opposite edge. Hence, we can assume both conclusion (i) and conclusion (ii) hold before we begin the normalization procedure.

We claim that before the normalization procedure, the total number of arcs of F° in shaded faces is

$$2\sum_i \#(F \cap D_i) = \sum_i \#(F \cap C_i).$$

This is because each component of $F^\circ \cap D_i$ runs from C_i to C_i and consists of one arc in each of the two shaded faces comprising D_i. Each such arc runs from an ideal vertex at C_i to the opposite edge, as in Figure 7.

Now, consider what happens during the normalization procedure of Lemma 3.4. That procedure monotonically reduces the complexity $c(F)$. In other words, for every face σ, $\#(F \cap \sigma)$ either stays constant or goes down. But by Corollary 2.2, the quantity $2\sum_i \#(F \cap D_i)$ is already minimal before normalization. Since this quantity is the

total number of intersections between F° and the shaded faces, it follows that $\#(F \cap \sigma)$ stays constant for every shaded face σ. This means that the intersections between F and the shaded faces remain as in Figure 7, and conclusions (i) and (ii) remain true throughout the normalization process. □

Lemma 3.6 *Assume that F° is in normal form. For each cusp torus T_i corresponding to a crossing circle C_i, each component of $\partial F^\circ \cap T_i$ consists of $(n_i - 1)$ segments parallel to shaded faces and 2 diagonal segments that have one endpoint on a white face and one endpoint on a shaded face. Here, n_i is the number of crossings in the twist region of C_i. See Figure 8.*

Proof Recall from [7, Lemma 2.6] that the cusp torus T_i corresponding to a crossing circle C_i is cut by the polyhedra into two rectangular boundary faces, one in each polyhedron. In the universal cover of T_i, we have a rectangular lattice spanned by $\boldsymbol{s}$ and $\boldsymbol{w}$, where $\boldsymbol{s}$ is a step parallel to a shaded face (horizontal in Figure 8) and $\boldsymbol{w}$ is a step parallel to a white face (vertical in Figure 8). In order to recover $S^3 \smallsetminus K$ from $S^3 \smallsetminus L$, we need to fill the torus T_i along a slope corresponding to the meridian of C_i in $S^3 \smallsetminus J$. By [7, Theorem 2.7], this Dehn filling slope is homologous to $\boldsymbol{w} + n_i \boldsymbol{s}$.

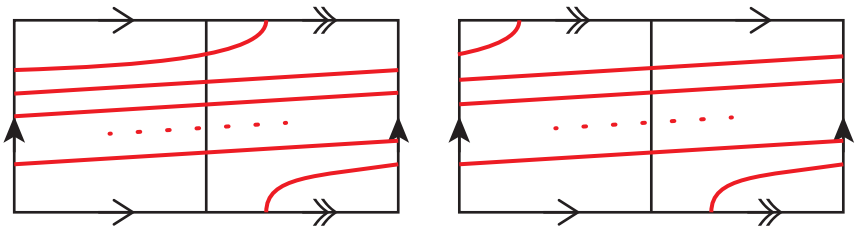

Figure 8: The cusp torus T_i of a crossing circle C_i is subdivided into two boundary rectangles. There are two combinatorial possibilities, depending on whether the number of crossings n_i in the twist region is even (shown on the left) or odd (shown on the right). The normal curve in T_i representing a component of ∂F° must be as shown in red.

In $S^3 \smallsetminus J$, the punctured surface F° meets the neighborhood of each crossing circle in a meridian. By the above paragraph, each component of ∂F° on T_i has homological intersection ± 1 with the shaded faces. On the other hand, by Lemma 3.5, each puncture of F° at C_i gives rise to a single arc in the shaded disk. Thus, each curve of ∂F° on T_i only has *geometric* intersection number 1 with the shaded faces. The only way to achieve this while staying in the homology class $\boldsymbol{w} + n_i \boldsymbol{s}$ is to take $n_i - 1$ segments parallel to $\boldsymbol{s}$ along with two diagonal segments whose sum is $\boldsymbol{w} + \boldsymbol{s}$. □

In the following section we need a vocabulary that allows us to translate combinatorial statements regarding normal loops in the boundary of P into combinatorial statements regarding the knot diagram $D(K)$. The following remark provides this translation.

Remark 3.7 The homeomorphism from $S^3 \smallsetminus J$ to $S^3 \smallsetminus L$ can be taken to be the identity outside a neighborhood of the union of the crossing disks in S^3. We can view this fact diagrammatically by shrinking the twist regions in the diagram of K until each is contained in a regular neighborhood of the arc of intersection between the corresponding crossing disk and the plane of projection for K. By Theorem 3.1, the white faces of the polyhedral decomposition of the complement of L meet the complement of the neighborhood of the union of the crossing disks in S^3 exactly in the plane of projection for L. Equivalently, the white faces of the polyhedral decomposition of the complement of J meet the complement of a regular neighborhood of the twist regions of K exactly in the plane of projection for K. In this way, arcs and loops in the white faces of the polyhedral decomposition are arcs and loops in the complement of the twist regions in the plane of projection for K.

Additionally, in light of Lemma 3.6, we know exactly how a normal surface meets the faces of the boundary of a polyhedron that correspond to cusp tori. In particular, if a normal loop in the boundary of a polyhedron meets only white faces and cusp tori faces, then each component of intersection with the cusp tori faces is a segment in the $\boldsymbol{s}$ direction. Hence, if a normal surface meets the boundary of a polyhedron in a loop that is disjoint from the shaded faces and this loop meets the collection of cusp tori faces in n components, then there is a curve in the plane of projection for K that cuts through twist regions n times and meets K in exactly $2n$ points. See Figure 9.

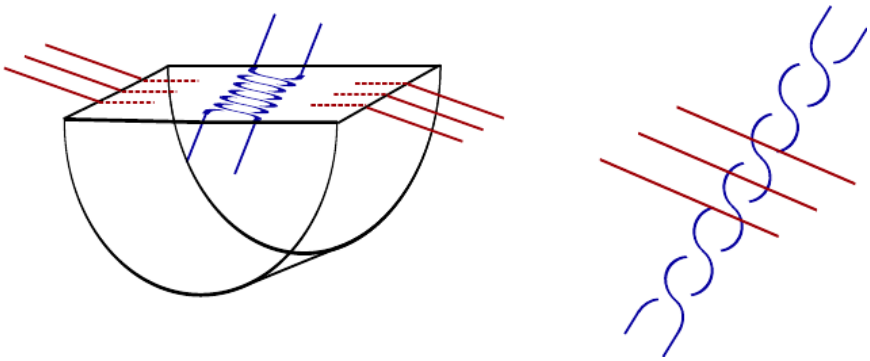

Figure 9: Curves in ∂P that are contained in the union of the white faces and the boundary faces give rise to curves in the plane of projection.

4 Intersections with crossing circles

In this section, we bound from below the number of times that a c–incompressible surface F must meet the crossing circles. We note that some, but not all, of the subsequent lemmas carry the hypothesis that F is closed. This will allow us maximum flexibility in proving Theorems 1.1 and 1.2.

Lemma 4.1 *Suppose that $F \subset S^3 \smallsetminus K$ is a closed, c–incompressible surface. Then F must intersect a crossing circle.*

Proof Suppose that F is disjoint from every C_i. Then $F = F^\circ$, and by Lemma 3.5 we can assume F is normal and disjoint from the crossing disks. By Theorem 3.1, the shaded faces of $P \cup P'$ glue to form the crossing disks. Thus $F \cap (\partial P \cup \partial P')$ is entirely contained in the white faces.

Since P and P' are checkerboard colored, every side of every white face borders a shaded face. But F is disjoint from the shaded faces, hence it cannot meet any edge of the white faces. Thus any intersection of F with a white face must be a simple closed curve, contradicting the normality of $F = F^\circ$. □

Lemma 4.2 *Suppose that $F \subset S^3 \smallsetminus K$ is a c–incompressible surface, either meridional or closed. Let $\Delta \subset F^\circ$ be a normal disk that meets exactly one crossing circle cusp. Then Δ must also meet a cusp corresponding to K.*

In particular, if F is closed, some normal disk $\Delta \subset F^\circ$ must meet at least two crossing circles.

Proof Let $\gamma \subset \partial\Delta$ be the unique arc of $\partial\Delta$ in a boundary face T_i corresponding to a crossing circle C_i. By Lemma 3.6, γ is either a segment in the $\boldsymbol{s}$ direction, parallel to a shaded face, or else a diagonal segment that runs from a white face to a shaded face. We consider these possibilities in turn.

Case 1 Suppose that γ runs parallel to the shaded faces, from a white face ω' to another white face ω. Consider where $\partial\Delta$ can go next. If $\partial\Delta$ crosses an edge of a polyhedron into a shaded face σ, Lemma 3.5 implies that it must next run into a boundary face T_j corresponding to some crossing circle C_j. But by hypothesis, $\partial\Delta$ meets only one boundary face, hence $T_i = T_j$. Thus, $\partial\Delta$ must meet T_i both in a segment parallel to a shaded face and in a diagonal segment that runs from a white face to a shaded face, contradicting normality.

If $\partial\Delta$ runs from ω directly into the boundary face T_j of a crossing circle C_j, then again we must have $T_i = T_j$, which means that $\omega = \omega'$ and $\partial\Delta$ contains only two

segments. But then, as Figure 10 shows, we can use Remark 3.7 to find a loop in $D(K)$ corresponding to γ that intersects K twice with nontrivial regions on each side. This contradicts the primeness of the diagram $D(K)$.

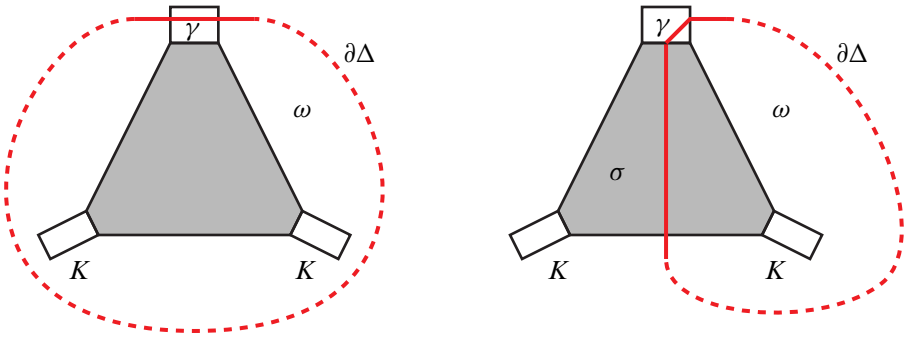

Figure 10: The two cases in the proof of Lemma 4.2. Left: $\partial\Delta$ must intersect a cusp of K, or else the dotted arc lies in a single white face, violating primeness. Right: $\partial\Delta$ must intersect a cusp of K, because the two endpoints of the dotted arc are separated by an odd number of knot strands.

The remaining possibility, if γ is a segment in the $\boldsymbol{s}$ direction, is that $\partial\Delta$ runs through ω to a truncated ideal vertex corresponding to K. This is our desired conclusion.

Case 2 Suppose now that γ is a diagonal segment that runs from a shaded face σ to a white face ω. Then, observe that the two ends of γ are separated by an odd number of knot strands. (See Figure 10.) Thus, to form a closed curve, $\partial\Delta$ must either cross a strand of K, which is our desired conclusion, or cross through another shaded face σ'. But then, as above, $\partial\Delta$ would have to run through σ' to a boundary face T_j of some crossing circle C_j, which contradicts either normality (if $T_i = T_j$) or the hypotheses (if $T_i \neq T_j$).

Thus, in all cases, $\partial\Delta$ must meet a cusp corresponding to K. □

Lemma 4.3 *Suppose that $F \subset S^3 \smallsetminus K$ is a closed, c–incompressible surface. Then F must intersect at least 3 crossing circles.*

Proof Suppose F meets strictly fewer than three crossing circles. By Lemmas 4.1 and 4.2, F must meet exactly two distinct crossing circles, C_i and C_j.

Recall that Lemma 3.6 implies that each component of intersection between ∂F° and the cusp torus T_i contains $n_i - 1 \geq 5$ segments parallel to $\boldsymbol{s}$ and only two other segments (and similarly for T_j). By Lemma 4.2, a normal disk $\Delta \subset F$ cannot intersect T_i only. It follows that some disk $\Delta \subset F$ intersects each of T_i and T_j in a segment parallel to $\boldsymbol{s}$.

Consider how the curve $\partial\Delta$ can close up. This curve cannot meet the ideal vertices corresponding to K, and it cannot meet any additional shaded faces either (otherwise Lemma 3.5 would force Δ to run into an additional crossing circle C_k). The only remaining possibility is that $\partial\Delta$ runs through a white face from T_i to T_j, and then through another white face back to T_i. Now, Figure 11 shows that we can use Remark 3.7 to find a loop in the projection plane corresponding to $\partial\Delta$ that intersects K four times, with two intersections adjacent to the twist region of C_i and the remaining intersections adjacent to the twist region of C_j. This violates the hypothesis that $D(K)$ is twist reduced. □

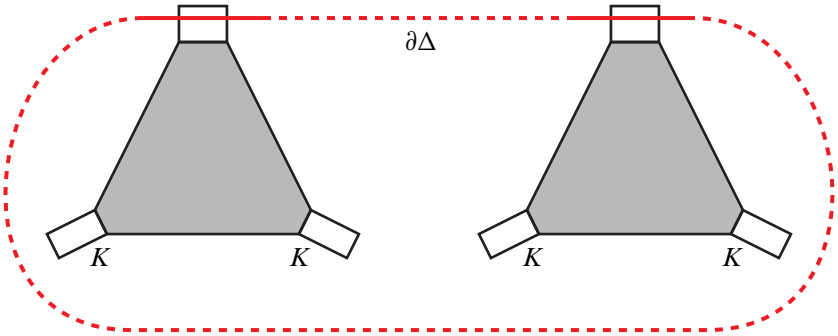

Figure 11: Lemma 4.3. If $\partial\Delta$ meets exactly two crossing circle cusps in segments parallel to $\boldsymbol{s}$, then the diagram $D(K)$ cannot be twist reduced.

Lemma 4.4 *Every cusp T_i of a crossing circle C_i contains an even number of components of ∂F°. Furthermore, if K is a knot and F is a closed surface, then every cusp T_i met by F° contains at least 4 components of ∂F°.*

Proof The first result is an immediate consequence of the fact that F is separating.

For the second conclusion, suppose that F is a closed surface and K is a knot. Since F is closed, it must separate S^3 into two components. The knot K must lie in one of these components. But every arc of $F \cap D_i$ separates the two strands of K that puncture the disk D_i. Therefore, since K lies on one side of F, the arcs of $F \cap D_i$ must come in pairs. Hence, if F intersects a crossing circle C_i at all, it must meet it at least 4 times. □

Corollary 4.5 *If F is closed, the punctured surface F° must meet the crossing circle cusps b times, where $b \geq 12$ if K is a knot and $b \geq 6$ if K is a link.*

Proof Immediate from Lemmas 4.3 and 4.4. □

5 Combinatorial length

The lemmas in the previous section give us a lot of control over the number of times that F° meets the cusps T_i corresponding to the crossing circles. To prove Theorems 1.1 and 1.2, we need to show that each component of $\partial F^\circ \cap T_i$ makes a substantial contribution to the Euler characteristic of F°, hence to that of F as well.

This can be done in one of two ways: either by estimating the *geometric length* of each component of ∂F° on a maximal cusp corresponding to the crossing circle C_i, or by estimating its *combinatorial length* in the sense of bounding the complexity of the normal disks comprising F°. The paper of Futer and Purcell contains readily applicable estimates on both combinatorial length and geometric length [7, Theorem 3.10 and Proposition 5.13], and either result would suffice for Theorem 1.1. We choose to pursue the combinatorial approach, because this is the approach that will generalize to meridional surfaces in Theorem 1.2.

The notion of combinatorial length was developed by Lackenby, as part of his study of Dehn surgeries on alternating knots [10]. The main idea is that the Euler characteristic of a surface can be controlled by understanding the intersections between that surface and the truncated ideal vertices in an ideal polyhedral decomposition.

Let us recap the key definitions. Every normal disk $\Delta \subset F^\circ$ has a well-defined *combinatorial area*, computed using the dihedral angles of the polyhedra in a manner that mimics the area formula for hyperbolic polygons.

Definition 5.1 Let D be a normal disk in a right-angled ideal polyhedron P, with the boundary faces of P lying on ∂M. Let n be the number of interior edges of P crossed by ∂D. Then the *combinatorial area* of D is defined to be

$$\operatorname{area}(D) = \frac{\pi}{2} n + \pi |\partial D \cap \partial M| - 2\pi.$$

Furthermore, the combinatorial area of a normal surface H is defined to be the sum of the combinatorial areas of all of its constituent normal disks and is denoted by $\operatorname{area}(H)$.

Proposition 5.2 (Gauss–Bonnet theorem) *Let $H \subset M$ be a normal surface in a 3–manifold with a right-angled polyhedral decomposition. Then*

$$\operatorname{area}(H) = -2\pi \chi(H).$$

Specializing to the case where $M = S^3 \smallsetminus L$, we have a way to "see" combinatorial area from the crossing circles.

Definition 5.3 Let Δ be a normal disk with respect to the polyhedral decomposition of $S^3 \smallsetminus L$. Let $\gamma_1, \dots, \gamma_n$ be the segments of $\partial\Delta$ that lie in boundary faces corresponding to crossing circles, and suppose that $n \geq 1$. Then, for each γ_i, we define

$$\ell(\gamma_i, \Delta) = \frac{\text{area}(\Delta)}{n}.$$

In other words, the area of Δ is distributed evenly among its intersections with the crossing circle cusps.

It is worth remarking that our definition of $\ell(\gamma_i, \Delta)$ differs slightly from the corresponding definition in Futer and Purcell [7, Definition 4.9]. The difference is that the latter definition divides the area of Δ among *all* segments of Δ in boundary faces, not just those corresponding to crossing circles. Definition 5.3 is designed to give stronger versions of some of the following estimates.

Lemma 5.4 *Let S be any normal surface in the polyhedral decomposition of $S^3 \smallsetminus L$. Then*

$$\text{area}(S) \geq \sum_i \ell(\gamma_i, \Delta),$$

where the sum is taken over all normal disks $\Delta \subset S$ and all segments of $\partial\Delta$ in crossing circle cusps.

Proof This is immediate, since Definition 5.3 ensures that the area of each disk is counted with the appropriate weight. Note that the inequality might be strict, because there may be normal disks in S that have positive area but do not meet any crossing circle cusps. □

In the case where S is the meridional, c–incompressible surface F°, we have a lot of control over the areas of disks and the corresponding combinatorial areas.

Proposition 5.5 [7, Proposition 5.3] *Let D be a normal disk in a polyhedron P of a right-angled polyhedral decomposition of M, such that ∂D passes through at least one boundary face. Let $m = |\partial D \cap \partial M|$. If D is not a bigon or an ideal triangle, then*

$$\text{area}(D) \geq \frac{m\pi}{2}.$$

Lemma 5.6 *Let $\Delta \subset F^\circ$ be a normal disk that meets n crossing circle cusps, where $n \geq 1$. Then, for each segment γ_i of $\partial\Delta$ in a crossing circle cusp,*

$$\ell(\gamma_i, \Delta) \geq \max\left\{\frac{\pi}{n}, \frac{\pi}{3}\right\}.$$

Proof Let m be the number of segments of Δ in *all* boundary faces (belonging either to K or to a crossing circle). By Lemma 4.2, $m \geq 2$. Furthermore, by Lemma 3.5, Δ cannot be a bigon (because the boundary of a bigon runs between two consecutive ideal vertices). If Δ is an ideal triangle, then $\text{area}(\Delta) = \pi$. Thus, by Proposition 5.5, $\text{area}(\Delta) \geq \pi$ in all cases. By Definition 5.3, it follows that

$$\ell(\gamma_i, \Delta) \geq \frac{\pi}{n}.$$

It remains to show that $\ell(\gamma_i, \Delta) \geq \pi/3$. If $n \leq 3$, we are done by the previous paragraph. Alternately, if $n > 3$, Proposition 5.5 gives $\text{area}(\Delta) \geq m\pi/2$, where $m \geq n$. Thus

$$\ell(\gamma_i, \Delta) = \frac{\text{area}(\Delta)}{n} \geq \frac{m\pi}{2n} \geq \frac{\pi}{2}. \qquad \square$$

We can now complete the proof of Theorem 1.1.

Proof of Theorem 1.1 Let $F \subset S^3 \smallsetminus K$ be a closed, c–incompressible surface. Isotope F into a position that minimizes the intersection number with the crossing disks D_i. After drilling out the crossing circles, we obtain a surface $F^\circ = F \cap S^3 \smallsetminus L$, which can be placed in normal form via the procedure of Lemma 3.4.

Let b be the number of boundary components of F° on the crossing circle cusps of $S^3 \smallsetminus L$. By Corollary 4.5, we have $b \geq 6$, with $b \geq 12$ in case K is a knot. Furthermore, by Lemma 3.6, each of these b components consists of $(n_i + 1)$ segments in boundary faces, where $n_i \geq h(D)$. Thus, by Lemma 5.6, each component of ∂F° contributes at least $(h(D) + 1)\pi/3$ to the area of F°.

Now, we may compute

$$\begin{aligned}
(1) \qquad -2\pi(\chi(F) - b) &= -2\pi\chi(F^\circ) && \text{by the construction of } F^\circ \\
&= \text{area}(F^\circ) && \text{by the Gauss–Bonnet formula} \\
&\geq \sum_i \ell(\gamma_i, \Delta) && \text{by Lemma 5.4} \\
&\geq \frac{\pi}{3} \cdot b \cdot (h(D) + 1) && \text{by Lemma 5.6,}
\end{aligned}$$

which simplifies to give

$$(2) \qquad \chi(F) \leq \frac{b}{6}(5 - h(D)).$$

Substituting $b \geq 6$ for links and $b \geq 12$ for knots gives the desired result. $\square$

The same ideas, with one added ingredient, also prove Theorem 1.2.

Proof of Theorem 1.2 Let $F \subset S^3 \smallsetminus K$ be a compact, connected, meridional, c–incompressible surface. Isotope F into a position that minimizes the intersection number with the crossing disks D_i. Drill out the crossing circles, and normalize F° in $S^3 \smallsetminus L$.

Unlike the setting of closed surfaces (that is, unlike Lemma 4.1), it may happen that F° is disjoint from the crossing circle cusps, ie $F^\circ = F$. Then, by Lemma 3.5, each normal disk $\Delta \subset F^\circ$ must be disjoint from the shaded faces. In other words, $\partial\Delta$ is a closed curve in the white projection plane, which intersects the cusps of K some number of times. This closed curve bounds a disk Δ in the polyhedron P, unique up to isotopy. Recall that P is glued to P' along all its white faces, and the gluing map is the identity on white faces. Thus we have an identical normal curve in P', which again bounds a normal disk Δ' that is unique up to isotopy. Since Δ and Δ' are glued to each other along all their edges, we conclude that $F = F^\circ$ is a sphere punctured some number of times by K and, by Remark 3.7, that it meets the projection plane for K along the single closed curve $\partial\Delta = \partial\Delta'$.

Next, consider what happens if ∂F° contains b components along the crossing circle cusps, where $b > 0$. By Lemma 4.4, b is even. If $b \geq 6$, then we argue exactly as in the proof of Theorem 1.1 and the computation (1) produces the same estimate (2) as for closed surfaces, namely

$$\chi(F) \leq 5 - h(D).$$

If $b = 2$, then ∂F° must intersect only one crossing circle C_i, and in particular every normal disk of F° has at most one segment along a crossing circle cusp. Thus Lemma 5.6 tells us that $\ell(\gamma, \Delta) \geq \pi$ for every segment γ along T_i, hence each component of ∂F° along T_i contributes at least $(h(D)+1)\pi$ to the area of F°. As a consequence, a calculation as in (1) gives

$$-2\pi(\chi(F) - b) \geq \pi \cdot b \cdot (h(D) + 1).$$

After substituting $b = 2$, this simplifies to

$$\chi(F) \leq 1 - h(D).$$

Similarly, if $b = 4$, then ∂F° intersects either one or two crossing circles. Consequently, every normal disk of F° has at most two segments along a crossing circle cusp. Thus Lemma 5.6 (with $n \leq 2$) tells us that $\ell(\gamma, \Delta) \geq \pi/2$ for every segment γ along a crossing circle cusp T_i. Hence, each component of ∂F° along T_i contributes at least $(h(D)+1)\pi/2$ to the area of F°, and a calculation as in (1) gives

$$-2\pi(\chi(F) - b) \geq \frac{\pi}{2} \cdot b \cdot (h(D) + 1).$$

After substituting $b = 4$, this simplifies to

$$\chi(F) \leq 3 - h(D). \qquad \square$$

Acknowledgements Futer is supported in part by NSF grants DMS-1007221 and DMS-1408682. Tomova is supported in part by NSF grant DMS-1054450.

References

[1] **C C Adams**, *Augmented alternating link complements are hyperbolic*, from: "Low-dimensional topology and Kleinian groups", (D B A Epstein, editor), London Math. Soc. Lecture Note Ser. 112, Cambridge Univ. Press (1986) 115–130 MR903861

[2] **D Bachman**, **S Schleimer**, *Distance and bridge position*, Pacific J. Math. 219 (2005) 221–235 MR2175113

[3] **R Blair**, **M Campisi**, **J Johnson**, **S A Taylor**, **M Tomova**, *Exceptional and cosmetic surgeries on knots* `arXiv:1209.0197`

[4] **R Crowell**, *Genus of alternating link types*, Ann. of Math. 69 (1959) 258–275 MR0099665

[5] **D Futer**, **F Guéritaud**, *Angled decompositions of arborescent link complements*, Proc. Lond. Math. Soc. 98 (2009) 325–364 MR2481951

[6] **D Futer**, **E Kalfagianni**, **J S Purcell**, *Dehn filling, volume, and the Jones polynomial*, J. Differential Geom. 78 (2008) 429–464 MR2396249

[7] **D Futer**, **J S Purcell**, *Links with no exceptional surgeries*, Comment. Math. Helv. 82 (2007) 629–664 MR2314056

[8] **D Futer**, **J S Purcell**, *Explicit Dehn filling and Heegaard splittings*, Comm. Anal. Geom. 21 (2013) 625–650 MR3078950

[9] **W Haken**, *Theorie der Normalflächen*, Acta Math. 105 (1961) 245–375 MR0141106

[10] **M Lackenby**, *Word hyperbolic Dehn surgery*, Invent. Math. 140 (2000) 243–282 MR1756996

[11] **M Lackenby**, *The volume of hyperbolic alternating link complements*, Proc. London Math. Soc. 88 (2004) 204–224 MR2018964

[12] **W Menasco**, *Closed incompressible surfaces in alternating knot and link complements*, Topology 23 (1984) 37–44 MR721450

[13] **K Murasugi**, *On the genus of the alternating knot, I and II*, J. Math. Soc. Japan 10 (1958) 94–105, 235–248 MR0099664

[14] **J S Purcell**, *An introduction to fully augmented links*, from: "Interactions between hyperbolic geometry, quantum topology and number theory", (A Champanerkar, O Dasbach, E Kalfagianni, I Kofman, W Neumann, N Stoltzfus, editors), Contemp. Math. 541, Amer. Math. Soc. (2011) 205–220 MR2796634

[15] **M Tomova**, *Multiple bridge surfaces restrict knot distance*, Algebr. Geom. Topol. 7 (2007) 957–1006 MR2336246

[16] **A Zupan**, *Bridge and pants complexities of knots*, J. Lond. Math. Soc. 87 (2013) 43–68 MR3022706

Department of Mathematics and Statistics, California State University, Long Beach
1250 Bellflower Blvd, Long Beach, CA 90840, USA

Department of Mathematics, Temple University
1805 North Broad St., Philadelphia, PA 19122, USA

Department of Mathematics, University of Iowa
14 MacLean Hall, Iowa City, IA 52242-1419, USA

ryan.blair@csulb.edu, dfuter@temple.edu, maggy-tomova@uiowa.edu

http://www.csulb.edu/~rblair/, http://www.math.temple.edu/~dfuter/, http://www.math.uiowa.edu/~mtomova/

Received: 12 February 2014 Revised: 4 August 2014